

# Quantitative diagnosis of continuous-valued, steady-state systems

Nicolas Rouquette

Jet Propulsion Laboratory

M/S 525-3950

Pasadena, California 91109

e-mail : Nicolas.Rouquette@jpl.nasa.gov

## Abstract

Quantitative diagnosis involves numerically estimating the values of unobservable parameters that best explain the observed parameter values. We consider quantitative diagnosis for continuous, lumped-parameter, steady-state physical systems because such models are easy to construct and the diagnosis problem is considerably simpler than that for corresponding dynamic models. To further tackle the difficulties of numerically inverting a simulation model to compute a diagnosis, we propose to decompose a physical system model in terms of feedback loops. This decomposition reduces the dimension of the problem and consequently decreases the diagnosis search space. We illustrate this approach on a model of a thermal control system studied in earlier research.

## Introduction

In a model-based reasoning setting, diagnosis is the problem of explaining differences between the behavior observed of a physical system and the behavior predicted from a model of that system. Regardless of the nature of the physical system model (logical, qualitative, quantitative, probabilistic, etc. ..), diagnosis is an example of an ill-posed inverse problem: Taken in isolation, behavior observations typically have a large number of possible explanations. Narrowing the set of likely explanations, also called hypotheses, to a manageable size is one of the fundamental problems of diagnosis [Davis and Hamscher, 1992].

In previous work on analog diagnosis, researchers have used various approaches to focus hypothesis generation: de Kleer and Brown [de Kleer and Brown, 1992] rely on a component-based ontology to map hypotheses onto individual components while others such as [Dvorak and Kuipers, 1992, deCoste, 1992, Oyeleye *et al.*, 1992] rely on the dynamic behavior of the system for diagnostic clues. Eventually, some hypotheses have to be tested. For steady-state models, this means that a model prediction will have to be computed and matched against sensor measurements from the physical system. Since diagnosis involves inferring the state of non-measurable properties, it means that a fault hypothesis corresponds to a set of exogenous

model parameter values. Testing a fault hypothesis then means verifying that the exogenous values make model predictions consistent with the anomalous sensor measurements.

As de Kleer and Brown observed in [de Kleer and Brown, 1992], this form of hypothesis testing is computationally expensive because the actual exogenous parameter values that best explain the anomalous sensor measurements are, after all, unknown. Furthermore, the number of parameter values observed is typically much smaller than the total number of parameters in the system model. The number of unknown parameter values thus far exceeds the number of known values due to observations. Although there has been several advances in diagnosing continuous physical systems with feedback at a qualitative level [Rose and Kramer, 1991] the general problem is quite hard as noted in [Biswas and Yu, 1993, Srinivas, 1994].

Here, we focus on the problem of computing exogenous parameter values from observed measurements. That is, we assume that adequate diagnostic hypothesis have already been selected and that some agency or program focused the attention of the diagnoser to determine the diagnostic values of specific exogenous parameters. To illustrate our discussion, we will use a simplified, steady-state, continuous, lumped-parameter model of an external-active thermal control system (EATCS) shown in Fig. 1, a two-phase ammonia thermal controller designed at McDonnell Douglas once considered for Space Station Freedom. This system is designed to transfer heat from the crew cabin and electronic equipment and radiate it in space. The vehicle for this transfer is two-phase ammonia at or below saturation temperature. Under normal circumstances, the thermodynamic cycle works as follows: at the evaporators, the heat load vaporizes liquid ammonia thereby producing a two-phase mixture. The rotary fluid management device separates vapor and liquid by centrifugation. The vapor is circulated to the condenser side where it is liquified. To ensure proper operation over a wide range of heat load conditions, the back-pressure regulating valve ensures that there is enough liquid ammonia flow to handle the heat load

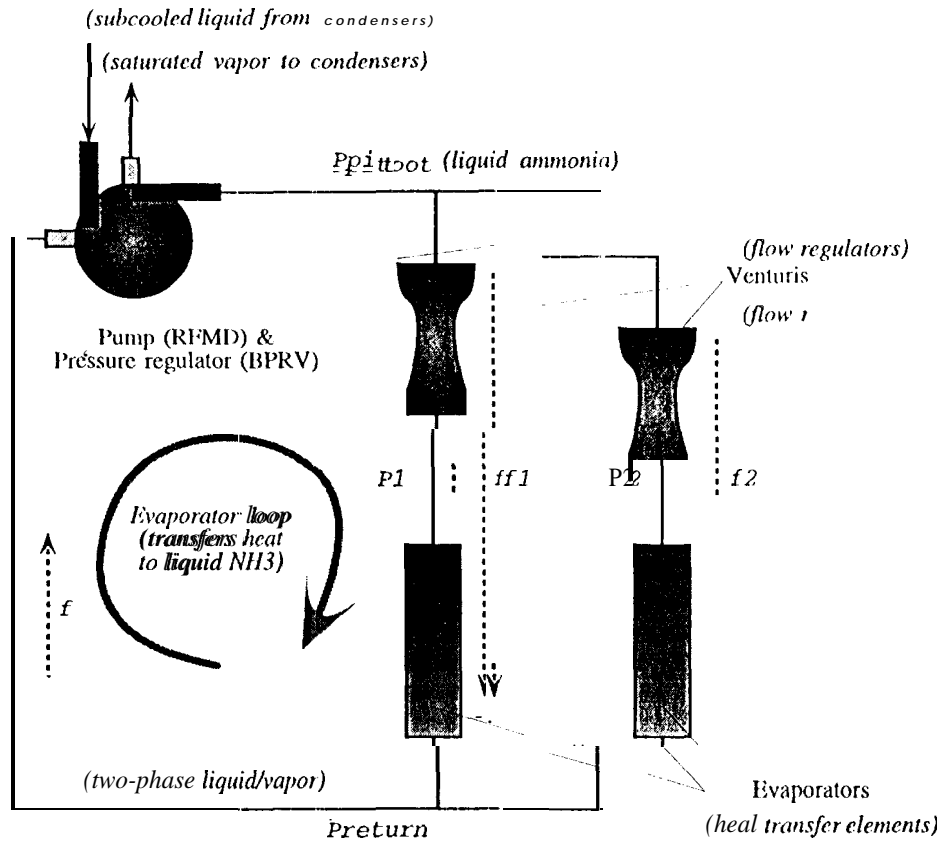


Figure 1: The evaporator loop of the EATCS.

without reaching a superheated vapor state'

We will first describe diagnosis as an inverse problem and then describe a technique for reducing the dimensionality of the diagnostic problem. Finally, we present an experiment to quantitatively evaluate this approach on simulated data.

### Diagnosis as an inverse problem

Let  $F(u, x) = 0$  stand for the equation model of the physical system where  $u$  is a set of exogenous parameters and  $x$  is the internal state of the system. Let  $y = G(x)$  be the observation function where  $y$  is the set of internal state parameters made observable by sensors attached to the physical system. Steady-state simulation is defined by the following problem:

Given  $u$  find  $y$  such that:

$$\begin{cases} F(u, x) = 0 \\ y = G(x) \end{cases}$$

Steady-state diagnosis is defined by the following inverse problem:

<sup>1</sup>Detailed analyses of this model appear in [Biswas et al., 1993, Rouquette, 1995]

Given  $y$  find  $u$  such that:

$$y = G(x) \quad (1)$$

where  $x$  is implicitly defined by:

$$F(u, x) = 0 \quad (2)$$

By combining Eq. 1 and Eq. 2 into one system of equations, this problem becomes equivalent to:

Given  $y$  find  $u$  such that:

$$S(u, y) = 0. \quad (3)$$

This problem is typically difficult because the number of exogenous parameters in  $u$  is not necessarily related to the number of observable parameters in  $y$ . In deed, the degree of sensor instrumentation of the physical system is often a separate consideration based on various factors such as cost, reliability, diagnosability, and available telecommunications bandwidth. These considerations typically lead to situations where the number of measured parameters, (i.e., the size of  $y$ ) is less than the number of exogenous parameters affecting the physical system (i.e., the size of  $u$ ).

Numerically, such situations are the most troublesome because the system  $S$  of equations is under-constrained and therefore may have an entire manifold

of solutions! Since a direct equation solving approach is not feasible, we need to compute a solution by minimizing the residual error of Eq. 3, i.e.,:

Given  $y$  and an error level  $\epsilon > 0$ , find  $u$  such that:

$$\|S(u, y)\| < \epsilon. \quad (4)$$

Recall that  $S$  implies that the entire set of state parameters,  $x$ , must be computed in order to solve the problem. Since there are far more internal parameters in  $x$  and  $u$  than we can typically observe in  $y$ , the least-squares minimization problem becomes numerically ill-posed, and there exists a manifold of possible solutions.

These two problems - solution manifolds and lack of guidance - imply that we need to exploit additional domain knowledge to better constrain the minimization problem and obtain meaningful solutions. This is a classic situation of ill-posed minimization problems which occurs quite frequently in solving integral equations of the first kind. For such problems, paradoxically it is less important to obtain the most accurate solution than deriving meaningful constraints that characterize an adequate solution. The difficulty of exercising this tradeoff makes least-squares minimization problems numerically delicate and difficult to solve.

For the diagnosis of physical systems, additional solution constraints stem from two sources: 1) the temporal behavior of the physical system and 2) the feedback structure of the physical system to reduce the effective number of internal parameters to be estimated (i.e., the size of  $x$ ).

### Temporal phases of steady-state behavior

Instead of solving the diagnosis problem for a given sample of observations  $y(t)$  measured at time  $t$ , we can exploit *a-priori* knowledge of the fact that the behavior of the physical system over time is approximately continuous and differentiable between discrete events (e.g. flipping a switch, opening or closing a valve). This leads us to consider not just one telemetry point but a temporal sequence of  $N$  such points in a sliding window from an instant in the past:  $t_N$  to the current instant:  $t = t_1$ :

$$y(t_N), \dots, y(t_2), y(t_1 = t).$$

Suppose now that the physical system is at steady state in the interval:  $[t_N, t_1]$ . This means that all measurements taken in that interval must be equal:  $y(t_N) = y(t_{N-1}) = \dots = y(t_1)$  by definition of steady-state behavior. This is equivalent to  $N - 1$  constraints:

$$\|y(t_{i-1}) - y(t_i)\| = 0 \quad \text{for all } i. \quad (5)$$

Suppose now that the physical system is not at steady state but that there are small transient or periodic changes small enough that, at some level of approximation, the system could be considered at or near

steady state. Then, it follows that in some sense, consecutive states are near each other as well. We can express this with  $N - 1$  constraints adapted from Eq. 5:

$$\|y(t_{i-1}) - y(t_i)\| < \epsilon. \quad (G)$$

Small  $\epsilon$  means a physical system close to steady-state. Conversely, large  $\epsilon$  means a physical system away from steady-state. Small  $\epsilon$  may be used to detect quiescence, while large  $\epsilon$  may be used to detect reachable steady-states nearby the current transient state.

We would like then to exploit multiple telemetry points in the following manner:

Given  $y(t_N), \dots, y(t_1)$  and an error bound  $\epsilon > 0$ , find  $u(t_N), \dots, u(t_1)$  such that:

$$\begin{cases} \|S(u(t_N), y(t_N))\| < \epsilon, \\ \|S(u(t_{N-1}), y(t_{N-1}))\| < \epsilon, \\ \vdots \\ \|S(u(t_2), y(t_2))\| < \epsilon, \\ \|S(u(t_1), y(t_1))\| < \epsilon. \end{cases} \quad (7)$$

As stated, we lack constraints to express the fact that consecutive telemetry observations  $y(t_N) \dots y(t_1)$  are not independent but instead correspond to an actual slice of behavior history. This type of additional knowledge fits within the application scope of *regularization* in parameter space, a technique of numerical analysis designed for nonlinear, ill-posed inverse problems [Chavent and Kunish, 1993]. Here, we use Tikhonov regularization [Chavent and Kunish, 1994] as follows. Instead of solving the diagnosis problem 7 as a nonlinear least-squares minimization, namely:

Given  $y$ , find:

$$\arg \min_{u(t_N), \dots, u(t_1)} \left\| \begin{bmatrix} S(u(t_N), y(t_N)) \\ \vdots \\ S(u(t_1), y(t_1)) \end{bmatrix} \right\|, \quad (8)$$

we augment problem 8 with additional information from  $n$ -th order difference relations for capturing the continuity of the physical system behavior. In the case of second-order difference relations, the resulting problem is shown below:

Given  $y$ , find:  $\arg \min_{u(t_N), \dots, u(t_1)}$

$$\left\| \begin{bmatrix} S(u(t_N), y(t_N)) \\ \vdots \\ S(u(t_1), y(t_1)) \\ \epsilon [u(t_N) - 2u(t_{N-1}) + u(t_{N-2})] \\ \vdots \\ \epsilon [u(t_3) - 2u(t_2) + u(t_1)] \end{bmatrix} \right\|. \quad (9)$$

If  $\epsilon = 0$ , then problem 9 is equivalent to nonlinear least-squares minimization problem 8. The value of  $\epsilon$  controls the importance attributed to the second-order differences terms, namely:

$$\epsilon [u(t_{k-1}) - 2u(t_k) + u(t_{k+1})].$$

A large regularization coefficient  $\epsilon$  means that the regularization term:  $u(t_{k-1}) - 2u(t_k) + u(t_{k+1})$  must be small which in turn forces the solution values to be close to each other. A small regularization coefficient  $\epsilon$  means that the regularization term can be correspondingly larger which in turn allows more freedom to the solution values. In practice, it can be difficult to select an adequate regularization coefficient. Thus, although powerful, this technique requires some fine tuning to work properly. Note that during the minimization process, several estimates of the exogenous parameters  $u(t_N), \dots, u(t_1)$  will be made, and the corresponding observable state parameters  $y(t_N), \dots, y(t_1)$  will therefore have to be computed by simulation. Since simulation consists in solving the state equations of the model, the use of regularization in parameter space as a means of performing diagnosis relies heavily on efficient equation solving techniques.

### Dimension reduction

In the previous section, we characterized the diagnosis problem as that of estimating the values of exogenous parameters,  $u$ , given observations  $y$  by solving  $S(u, y) = 0$  for a given  $y$ . As we saw in equations 1 and 2, this necessitates to find the values of the unobservable internal parameters  $x$  such that  $y = G(x)$  where  $F(u, x) = 0$ . The more unobservable parameters there are, the more difficult will it be to numerically compute a diagnosis,  $u$  and the more uncertain will the results be. Like all inverse problems where it is not directly possible to compute  $u$  for a given  $y$ , the solution is first estimated,  $\hat{u}$ , and a resultant,  $\hat{y}$ , is computed. Other estimates of  $u$  successively generated until the residual error,  $\|y - \hat{y}\|$ , is below a convergence threshold.

For a given estimate,  $\hat{u}$ , it is therefore important to minimize the number of unobservable state parameters,  $x$  that must be also estimated in order to compute a prediction  $\hat{y}$ . We address this issue by imposing a structure to the physical system model such that the size of  $x$  is as small as possible. To achieve this, we seek to identify the structure of the physical system model to obtain the set of independent state parameters that are equivalent to  $x \cup u$ . Here, we approximate this ideal situation by analyzing the feedback structure of the model so that  $x \cup u$  becomes as small as possible. This process occurs in three phases.

First, we start by constructing an algebraic ordering of the model parameters and equations to capture how parameters can be under, properly or over-constrained from equations and how equations can be properly constraining one or over-constraining multiple parameters. Figure 2 shows two examples with equations (top), algebraic orderings (middle) and corresponding high-level equation-solving code (bottom). This notion of algebraic ordering bears close resemblance to that of

causal ordering; the latter seeks to identify which parameters causally influence the values of other parameters while the former seeks to identify how parameters are computed from other parameters and equations. The distinction between the two stems from the differences between the causal aspects of a model equations and parameters and the computational aspects of numerically computing parameter values.

Second, we use graph-theoretic techniques to decompose the algebraic ordering in terms of strongly-connected components and each connected component in terms of feedback loops. Although there are theoretical limitations on solving the feedback vertex set problem for arbitrary directed graphs, the structure of algebraic orderings are strongly biased to reflect either physical feedback loops or algebraic circularities of dependencies. We have exploited these biases to construct efficient decomposition algorithms described in [Rouquette, 1995]. For the EATCS system, the algebraic ordering constructed is shown in Fig. 3 and the structural feedback decomposition made of this ordering is shown in Fig. 4. An example of algebraic equation, eq20, is shown below:

```
(defeqn :cond # [P1<PsatPitot*Lambda1
+ Ppitot*(1-Lambda1)]#
:if-true # [ f1=(Pi/4)*(Phi1/12.0)-2
*sqrt(2*rLiqPitot*(Ppitot-PsatPitot)
*4633.05/Lambda1)]#
:if-false # [f2=(Pi/4)*(Phi2/12.0)^2
*sqrt(2*rLiqPitot*(Ppitot-P2)
*4633.05/Lambda2)]#)
```

The third and final step consists in generating a simulation program from a decomposed algebraic ordering. The additional structure brought by feedback analysis allows us to eliminate all but the state parameters of each feedback loop thereby leading to a drastic dimension reduction in the simulation program as shown in Table I for the EATCS model. This three-stage approach to constructing simulation programs is implemented in a computer program, DAGGER described in [Rouquette, 1995].

### Experiment Setup

Sensor technology creates inherent limitations on what system parameters can be measured. For example, while it is possible to measure pressures, temperatures, flow rates or pump speeds, it is not yet feasible to measure other important characteristics, such as the quality of two-phase fluids indicating the ratio of vapor to liquid present in a fluid (1.0 for 100% liquid and 0% vapor; 0.0 for 0% liquid and 100% vapor). For the EATCS, we focused on the measurable parameters of Table 2. Depending on the modeling assumptions used, the total number of parameters (i.e.,  $|u| + |x| + |y|$ ) varies between 45 and 88 which makes this inverse problem undeniably severely under-constrained.

The diagnosis problem consists in identifying the values of the exogenous parameters which lead to a

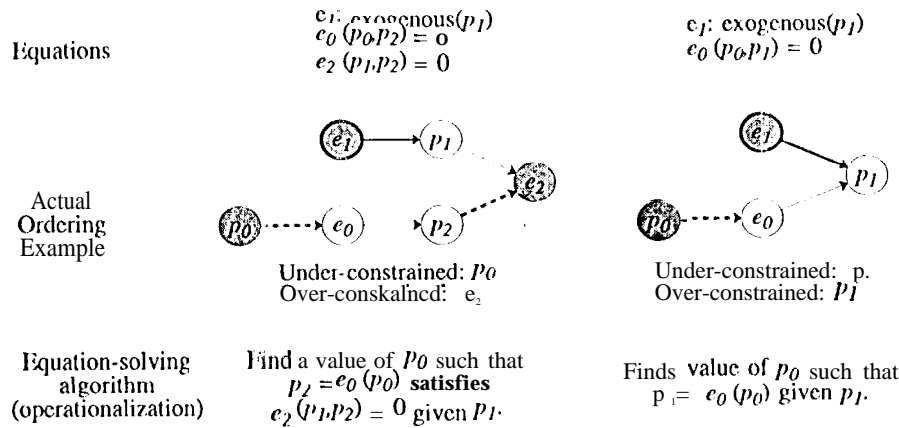


Figure 2: An example of over-constrained equation (left) and over-constrained parameter (right).

Loop	Number of parameters	
	Before decomposition	After decomposition
FLoop1	10	1
FLoop2	10	1
FLoop3	4	1

Table 1: Dimension reduction due to feedback decomposition

Sensor	Description
Pitot	Pitot outlet pressure.
Titot	Pitot outlet temperature.
P1	Pressure at first evaporator inlet.
P2	Pressure at second evaporator inlet.
T1	Temperature at first evaporator outlet.
T2	Temperature at second evaporator outlet.
TReturn	Temperature at BPRV inlet.
f	Flow rate at BPRV inlet.
f1	Flow rate at first venturi outlet.
f2	Flow rate at second venturi outlet.

Table 2: A set of measurable parameters of the EATCS

match between the simulated and observed values of the above 10 parameters. All the other internal state parameters (29 to 72 depending on the assumptions chosen) must be estimated as well. There are various ways to constrain this problem better for example, using ranges of physically plausible values for each estimated parameter and qualitative information about likely directions of change (e.g., [Kapadia *et al.*, 1994]). Eventually, these sources of information will narrow the problem down to intervals of plausible values in which a diagnosis must be found.

To evaluate the feasibility of doing quantitative diagnosis, we took a model of the EATCS and constructed two diagnosis engines. The first engine, the single-state diagnoser, attempts to identify the exogenous parameters that best explain a set of observations. The sec-

Sensor	Description
nMotor	Pump speed.
Preturn	Two-phase return pressure.
Phi1	Venturi 1 diameter.
Phi2	Venturi 2 diameter.
W1Power	Evaporator 1 load.
W2Power	Evaporator 2 load.

Table 3: A representative set physical characteristics of the EATCS

ond engine, the regularized diagnoser, exploits the near steady-state properties of a temporal sequence of observations to better constrain the problem.

**Single state dx** Given an observation vector  $y$ , find the internal parameters  $u$  which best explain the observation (Fig. 5).

**Regularized dx** Given  $N$  observations, use regularization in parameter space to find a sequence of  $N$  internal parameters which best explain the observations (Fig. 6).

The structures of the single-state and regularized diagnosis experiments are very similar (Figs. 5, 6). We generated telemetry data by first simulating states of the EATCS. This consists in describing a set of hypothetical physical circumstances in terms of values for the exogenous parameters (top left corner) and solving the EATCS model equations with respect to the exogenous parameters with a quantitative simulation

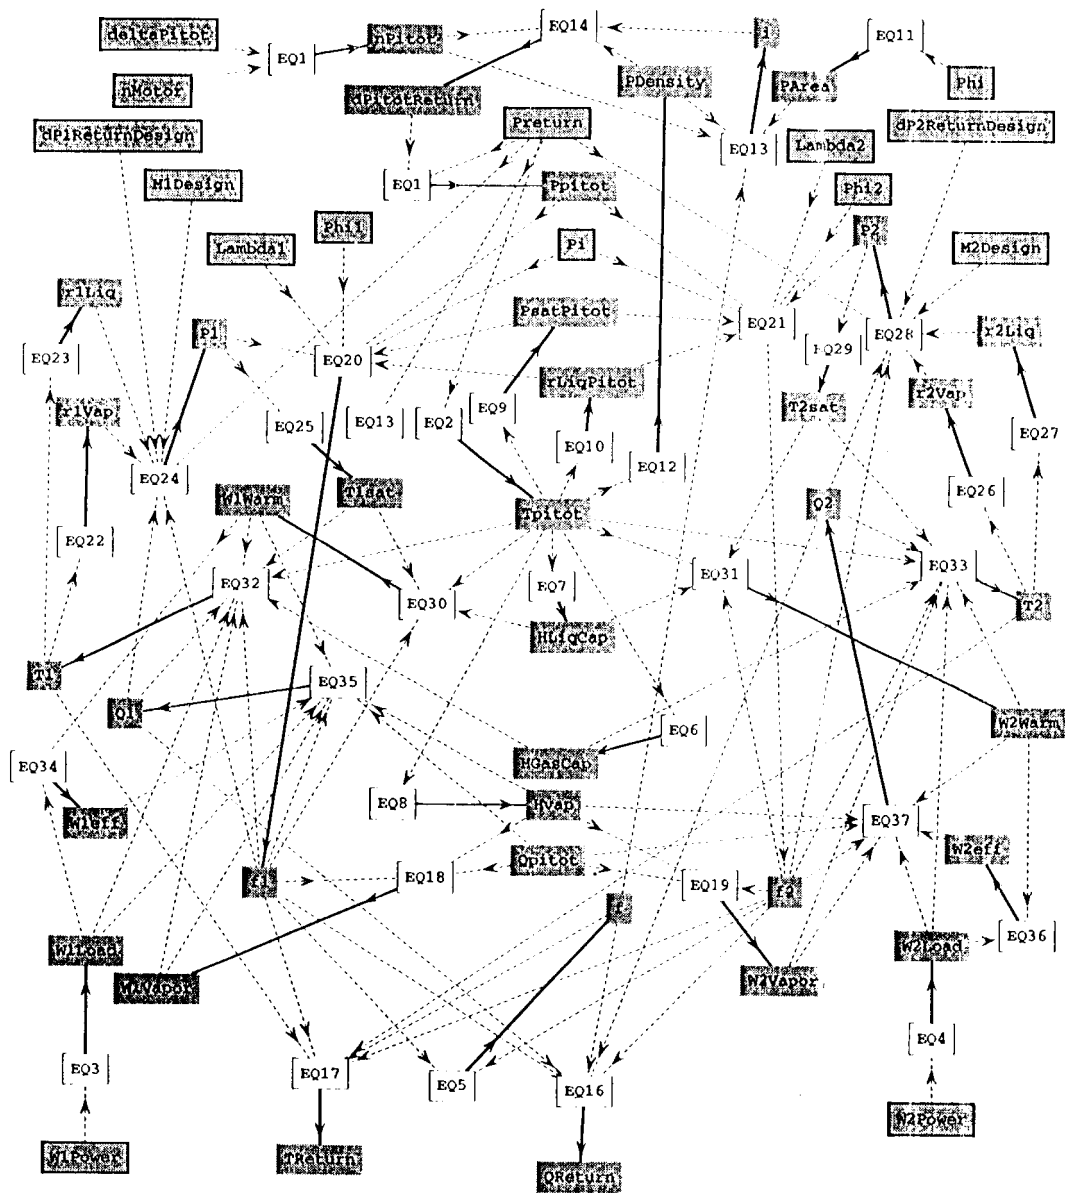


Figure 3: Algebraic Ordering of the FATES model. Solid edges represent how equations properly constrain parameters, dashed edges represent how parameters values are necessary to solving equations. Parameters are shown as gray boxes, exogenous parameters are emphasized with heavy lines; equations are shown with white boxes.

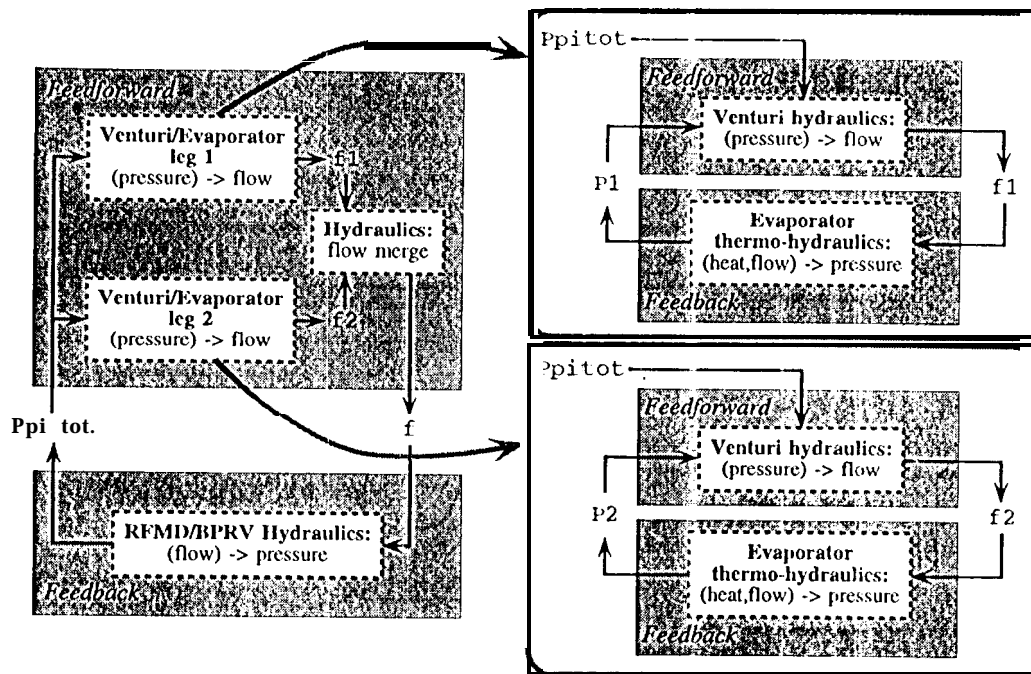


Figure 4: The 2-level hierarchical feedback decomposition of the EATCS model; gray boxes indicate the feedforward and feedback components of each feedback loop.

model.<sup>2</sup> At this stage, the values of the model parameters correspond to the steady-state configuration of the EATCS. We then selected a subset of the parameters to represent a set of plausible sensor measurements we could obtain from the actual EATCS hardware system. To these simulated values, we added either white or gaussian noise to account for various sources of noise such as sensor instruments, the physical process itself, and the approximative nature of the EATCS model used (top right corner).

The task of the diagnoser is to reconstruct the state of the physical system which when simulated, is the closest approximation of the true physical state of the EATCS (bottom). This works as follows: First, an estimate of the exogenous parameters generated (bottom right corner). Solving the model equations with respect to the exogenous parameter estimates then results in a set of predicted state parameter values. Comparing the actual telemetry observations with the predictions produces a set of residuals (center). The heart of the diagnosis then consists in adjusting the exogenous parameter value estimates until the residual prediction error is minimized (bottom left). Finally, the solution of the nonlinear least-squares problem defines a diagnosis whose accuracy is measured with respect to the original exogenous conditions (middle left).

<sup>2</sup>Due to the inability of conventional simulation methods to properly converge on physically plausible states, we were limited in the choice of simulation models to those constructed by DAGGER.

The structure of the regularized diagnoser is similar to the single-state except that the nonlinear least-squares minimization operates on the residual prediction error and a set of  $k$ -th order difference relations also known as regularized constraints.

The accuracy and performance of the simulation model are critical to the success of this approach. For this purpose, we used a graph-theoretic approach to analyzing steady-state, lumped-parameter models of algebraic equations in order to construct a specialized simulator whose structure espouses the feedback present in the model.

## Results

In most cases, the single-state diagnoser performed worse than the regularized diagnoser because there are typically many possible ways to explain a single observed state. Figure 7 shows a comparison of the single-state and the regularized diagnosers (1)x). Both diagnosers had to estimate the values of the following exogenous parameters, the pump speed (Fig. 7-a) (also directly measured through a sensor), the first venturi diameter (Fig. 7-b), and the heat load on the first evaporator (Fig. 7-c). Both diagnosers were given hydrothermal measurements for the first leg, namely, the evaporator outlet temperature (Fig. 7-d), the evaporator inlet pressure (Fig. 7-e) and the venturi flow rate (Fig. 7-f).

Although the regularized diagnoser clearly outperformed the single-state diagnoser, we were not able to

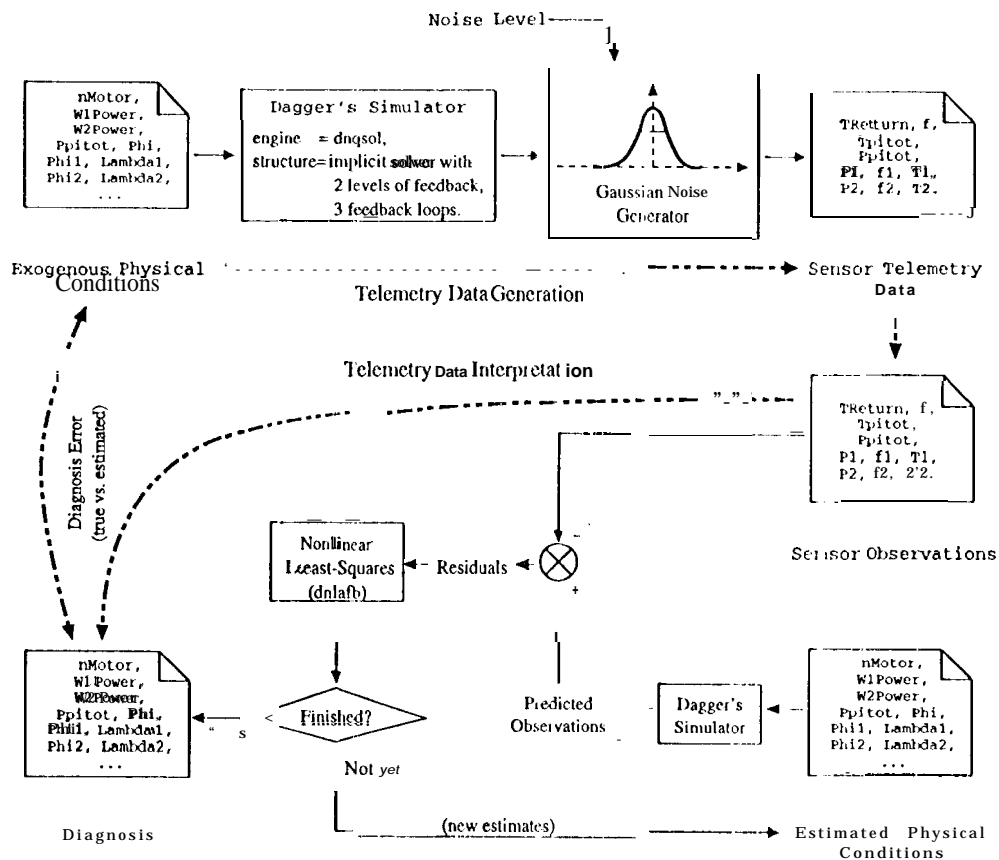


Figure 5: Single-state diagnosis as a numerical inverse problem: given  $n$  sensor observations, find a diagnosis as the values of the  $p$  non-observable internal state parameters that best explain the observations through steady-state simulation.

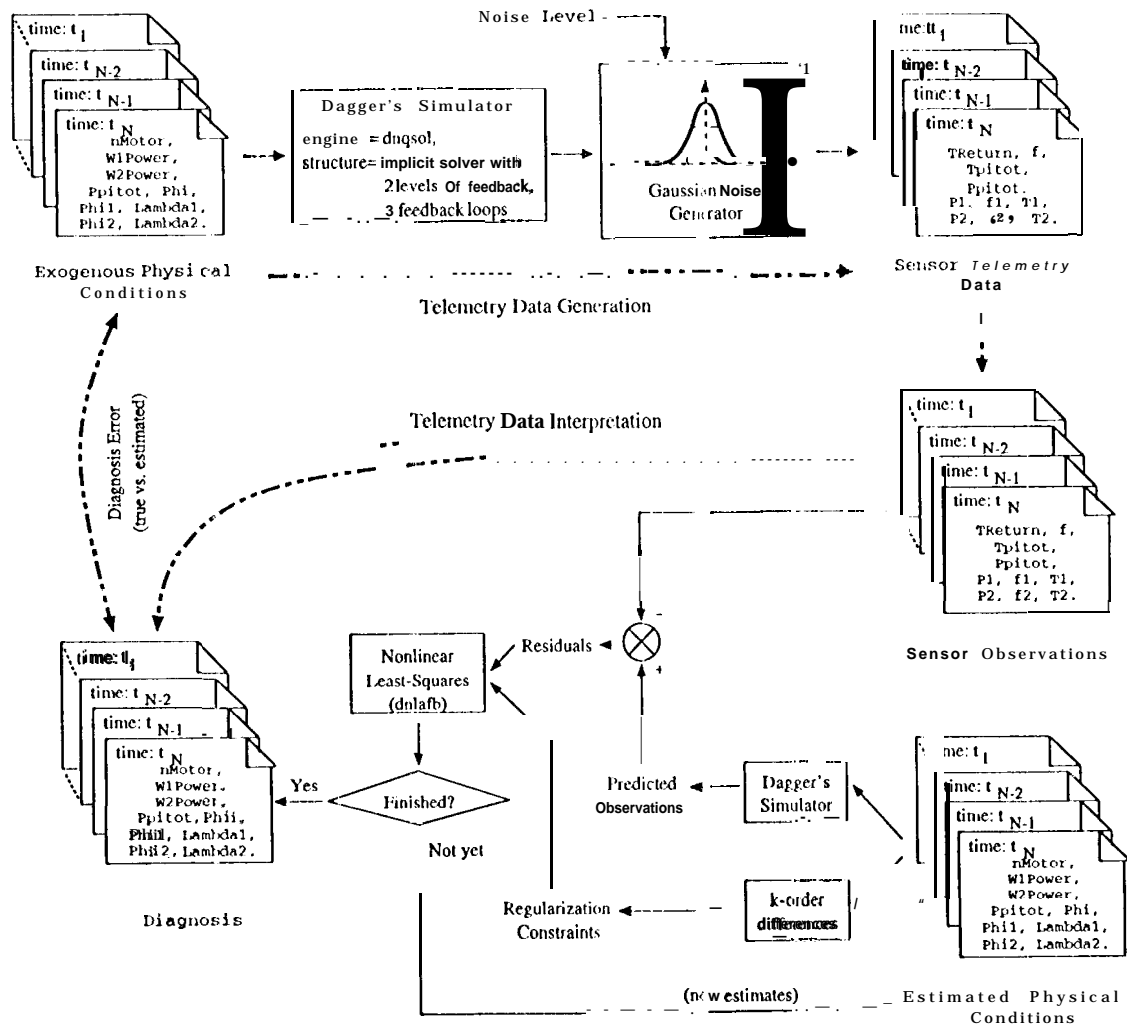


Figure 6: Regularized diagnosis is still a numerical inverse problem in which we further assume that an  $N$ -long temporal sequence of  $n$  sensor observations has a certain smoothness by virtue of each observation being near a steady state. By continuity near steady states, the internal state parameters leading to the temporal sequence of observations must have a corresponding degree of smoothness that we estimate in terms of  $k$ -order difference relations (the regularization constraints). A diagnosis thus becomes an  $N$ -long sequence of  $p$  non-observable internal state parameters that best explain the  $N$ -long temporal sequence of observations through Steady-state! Simulation.

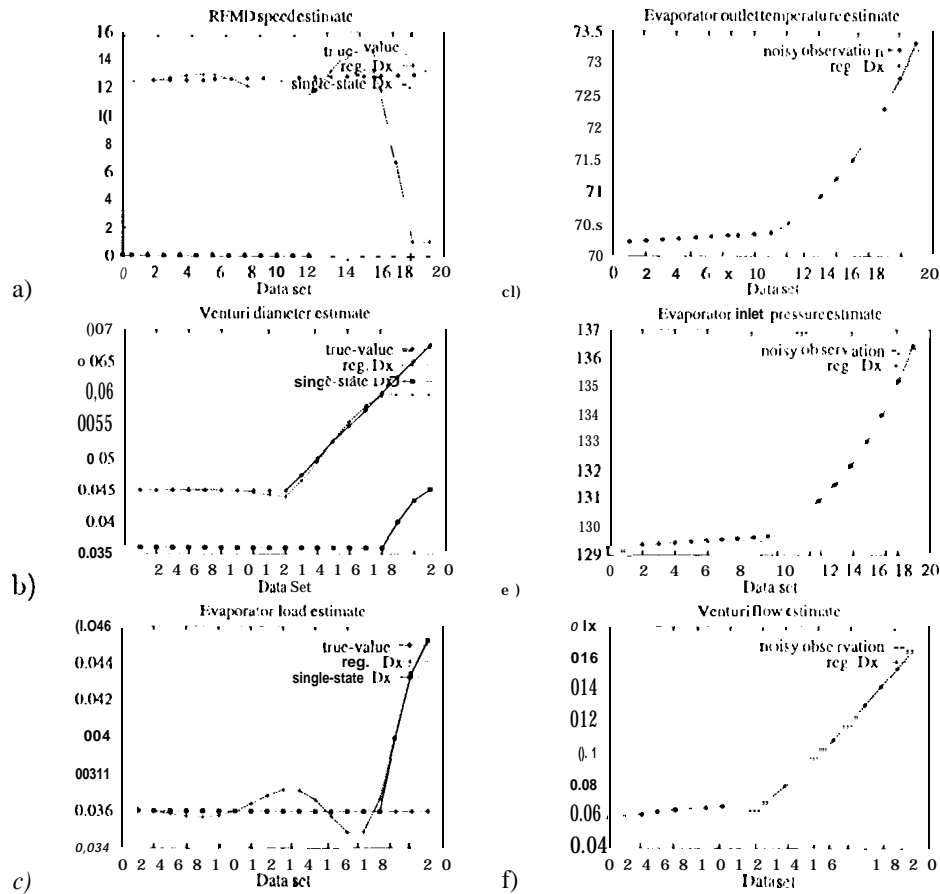


Figure 7: Noise-free diagnosis experiment for RFMD pump speedup, heat load increase and venturi clog. Although the predictions of temperature (cl), pressure (e), and flow rate (f) perfectly match the observations, there are significant differences between the true and diagnosed values of the exogenous parameters: RFMD speed (a), venturi diameter (b) and evaporator heat load (c). However, in all cases the regularized diagnoser was able to make reasonable estimates of these unknown values. On the other hand, the single-state diagnoser could not even produce a reasonable guess of the unknown values. This difference is explainable by the number of possible explanations there are of the behaviors seen in (d, e, and f): only multiple observations can help a system eliminate bad guesses and retain the good ones.

extend these results to noisy measurements. While the single-state diagnoser is definitely out of consideration for noisy diagnosis, we had originally expected to use regularization to impose global smoothness constraints that a diagnosis must exhibit. Unfortunately, instead of obtaining smooth diagnoses in the sense that the transitions from one state to the next are small, the regularized diagnoser overfitted the noise. The numerical analysis work required to properly tune a regularized diagnoser extended WCII beyond the scope of this thesis; therefore, we leave these investigations for future work.

## References

- Biswas, G. and Yu, X. 1993. A formal modeling scheme for continuous systems: Focus on diagnosis. In *Proceedings of the Eleventh National Conference on Artificial Intelligence*. The AAAI Press.
- Biswas, G.; Yu, X.; and Kapadia, R. 1993. Documentation: Diagnosis of thermal bus system for space station freedom. Technical report, Dept. of Computer Science, Vanderbilt University, Nashville, TN.
- Chavent, G. and Kunisch, K. 1994. Convergence of tikhonov regularization for constrained ill-posed inverse problems. *Inverse Problems* 10:63-76.
- Chavent, G. and Kunisch, K. 1993. Regularization in state space. *Mathematical Modeling and Numerical Analysis* 27(5):535-564.
- Davis, R. and Hamscher, W. 1992. Model-based reasoning: Troubleshooting. In Hamscher, W.; Console, L.; and de Kleer, J., editors 1992, *Readings in Model-Based Diagnosis*. Morgan Kaufmann, San Mateo, CA. 3-28.
- de Kleer, J. and Brown, J. S. 1992. Model-based diagnosis in sophie ii. In Hamscher, W.; Console, L.; and de Kleer, J., editors 1992, *Readings in Model-Based Diagnosis*. Morgan Kaufmann, San Mateo, CA. 179-205.
- deCoste, D. 1992. Dynamic across-time measurement interpretation. In Hamscher, W.; Console, L.; and de Kleer, J., editors 1992, *Readings in Model-Based Diagnosis*. Morgan Kaufmann, San Mateo, CA. 255-261.
- Dvorak, D. and Kuipers, B. 1992. Model-based monitoring of dynamic systems. In Hamscher, W.; Console, L.; and de Kleer, J., editors 1992, *Readings in Model-Based Diagnosis*. Morgan Kaufmann, San Mateo, CA. 249-254.
- Kapadia, R.; Biswas, G.; and Robertson, C. 1994. Dot: A framework for monitoring and diagnosis of continuous-valued systems. In Faltings, B. and Struss, P., editors 1994, *Fifth International Workshop on Principles of Diagnosis*, Fifth International Workshop on Principles of Diagnosis.
- Oyeleye, O. O.; Finch, F. E.; and Kramer, M. A. 1992. Qualitative modeling and fault diagnosis of dynamic processes by midas. In Hamscher, W.; Console, L.; and de Kleer, J., editors 1992, *Readings in Model-Based Diagnosis*. Morgan Kaufmann, San Mateo, CA. 262-275.
- Rose, P. and Kramer, M. A. 1991. Qualitative analysis of causal feedback. In *Proceedings of the Ninth National Conference on Artificial Intelligence*. The AAAI Press. 817-823.
- Rouquette, N. 1995. *Operationalizing Engineering Models of Steady-State Equations into Efficient Simulation Programs*. Ph.D. Dissertation, Dept. of Computer Science.
- Srinivas, S. 1994. A probabilistic approach to hierarchical model-based diagnosis. In Faltings, B. and Struss, P., editors 1994, *Fifth International Workshop on Principles of Diagnosis*, Fifth International Workshop on Principles of Diagnosis.